arXiv:1610.06549v1 [cs.CR] 20 Oct 2016

Untraceable VoIP Communication based on DC-nets

Christian Franck and Uli Sorger

Abstract

Untraceable communication is about hiding the identity of the sender or the recipient of a message. Currently most systems used in practice (e.g., TOR) rely on the principle that a message is routed via several relays to obfuscate its path through the network. However, as this increases the end-to-end latency it is not ideal for applications like Voice-over-IP (VoIP) communication, where participants will notice annoying delays if the data does not arrive fast enough.

We propose an approach based on the paradigm of Dining Cryptographer networks (DC-nets) that can be used to realize untraceable communication within small groups. The main features of our approach are low latency and resilience to packet-loss and fault packets sent by malicious players. We consider the special case of VoIP communication and propose techniques for a P2P implementation. We expose existing problems and sketch possible future large-scale systems composed of multiple groups.

1 Introduction

A few decades ago telecommunications were mainly realized using circuit switching, such that an electrical signal carrying the audio information could be sent from one correspondent to another. Nowadays it is more common that speech is digitized and then sent in small packets over a private network or over the internet.

In such real-time VoIP communication the quality of the user experience is dependent on the latency of the network. While some package-loss can be tolerated, the one-way end-to-end latency should ideally not exceed 150ms [1]. To this effect a stream of small packets is sent at a fast pace, for instance around 50 packets per second [5].

We are interested to have a system in which two users can communicate with each other via VoIP anonymously to the others. That means these two users know that they are speaking to each other, but nobody else can infer who is communicating with whom.

One approach would be to use an anonymisation system like the well-known TOR network [9], or alternatively one of the more recently proposed systems like Drac [8] or Herd [19]. However, in all these systems the security is based on the principle that the message is relayed several times on its way through the network, so that it is hard to distinguish who is the sender or the recipient of a message. One problem of this is that users have to trust the operators of intermediary relays and mixes that they are honest and do not disclose any information. Another problem is that this relaying increases the latency and the jitter, which reduces the quality of phone calls or video-conferences.

Another approach would be to use Dining Cryptographer Networks (DC-nets). Here, messages are not relayed but multiple players simultaneously send ciphertexts to the recipient, and then the recipient combines these ciphertexts to obtain a message. The recipient knows that one of the players must have sent the message, but he is not be able to distinguish which one. As opposed to relay based systems like TOR and mixing, the latency is lower, no central authority is required and might therefore be interesting for applications which require a low latency and for P2P scenarios.

However, in classical DC-nets the recipient can only recover the message once he has obtained all the ciphertexts. Furthermore, all these ciphertexts must be correct; if a malicious player provides faulty data the recipient cannot recover the message. Thus, if packets are late, lost or faulty, then no useful information is transmitted.

In this paper, we propose a protocol that is based on the paradigm of DC-nets but adapted for VoIP streaming. It enables two players out of a group of n players to communicate without disclosing to anybody else that they are communicating. Our protocol is resilient to packet-loss and to faulty packets, and can be implemented using lightweight cryptography. We study the performance, discuss practical imple-

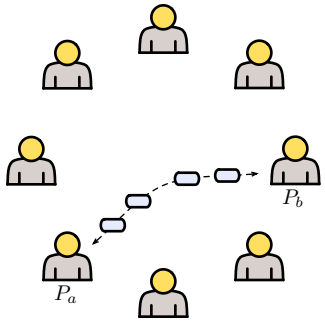

Figure 1: We want to enables two players P_a and P_b out of a group of n players $P_1, ..., P_n$ to communicate with VoIP in such a way that nobody except them can distinguish which pair of the $n \cdot (n-1)$ possible pairs of players is communicating.

mentations and propose possible future work.

The paper is organized as follows. In Section 2 we describe our model and our security requirements. In Section 3, we derive our new protocol. In Section 4, we discuss practical considerations, e.g. an implementation as a P2P protocol. In Section 5 we consider the performance. In Section 6 we sketch possible future work, and in Section 7 we review related work. In Section 8 we conclude with some remarks.

2 Model and Definitions

We consider the case where two players want to communicate and stay anonymous w.r.t. to the others, so our requirements are different from those of the original DC-net protocol. In the original DC protocol the recipient does not know the identity of the sender and the sender does not know the identity of the recipient, but in our model we assume that the two correspondent know each other's identity.

2.1 Model

We assume a communication network with n players $P_1, ..., P_n$ and an aggregator A, and we consider the problem where two of these players, P_a and P_b, want to communicate with each other such that nobody else can see that they are communicating.

After an appropriate setup, the messages are exchanged in successive transmission rounds which consist of two phases:

1. A collection phase where each player send data to the aggregator.
2. A broadcast phase during which the aggregator sends the aggregated data back to the participants.

This star topology is a well-known scheme for implementing a DC-net in networks without physical broadcast.

2.2 Security Properties

In this section we define the properties of correctness and privacy we want to achieve.

Correctness We say that a protocol is correct if, assuming the players P_a, P_b and the aggregator A participate correctly, it holds that:

- if the message of P_a encoding m_a reaches the aggregator A in time and the result forwarded by the aggregator A reaches P_b, then P_b can compute m_a; and
- if the message of P_b encoding m_b reaches the aggregator A in time and the result forwarded by the aggregator reaches P_a, then P_a can compute m_b.

Thus, if the players P_a, P_b and the aggregator A behave correctly and their packets arrive timely, then P_b and P_a can respectively recover m_a and m_b.

Privacy We say that a protocol is private if a computationally bounded adversary controlling up to $n-3$ players (other than P_a and P_b) and the aggregator A cannot determine who are the two correspondents with a probability better than random guessing.

For a set of honest participants $H \subset \{P_1, ..., P_n\}$ there are $c = |H| \cdot (|H|-1) \cdot 2^{-1}$ different pairs of correspondents. Thus a protocol is private if the adversary has a probability of guessing correctly the pair of correspondents which is not better than a random guess, i.e., with probability $1/c$.

3 A Protocol for Untraceable Streaming

To start we consider a simple protocol that meets the requirements from the previous section in an ideal situation and we progressively adapt it to end up with a protocol that can be used in a real-world scenario.

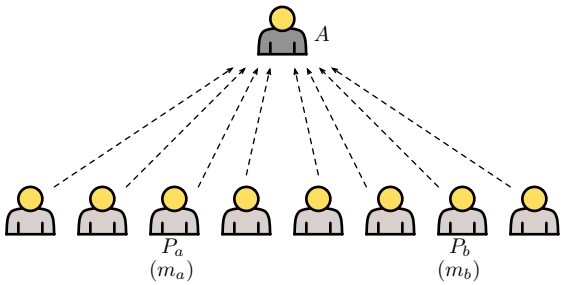

(a) Collection phase. Each player $P_1, ..., P_n$ sends a ciphertext to the aggregator A. The ciphertexts of players P_a and P_b encode respectively the messages m_a and m_b, whereas the other ciphertexts are 'empty'. The aggregator A aggregates the received ciphertexts.

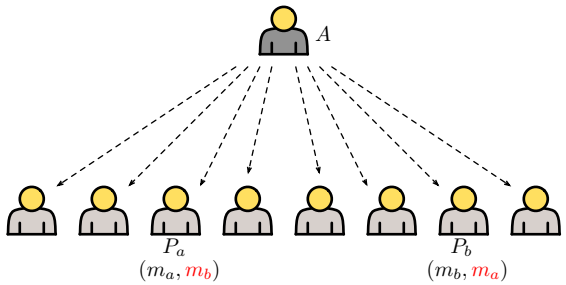

(b) Broadcast phase. The aggregator A sends the result of the aggregation back to all players $P_1, ..., P_n$. From this the players P_a and P_b can compute m_b and m_a.

Figure 2: We consider a system with n players $P_1, ..., P_n$ and an aggregator A, in which each transmission round comprises a sending and a broadcast phase.

Protocol 1 A simple protocol based on DC-nets.

1. **Setup**: $P_1, ..., P_n$ respectively have secret keys $k_1, ..., k_n$ such that $k_1 + ... + k_n = 0$. P_a additionally has m_a and P_b additionally has m_b.
2. **Collection phase**: Every P_i for $i \neq a, b$ sends $O_i = k_i$ to A. P_a sends $O_a = k_a + m_a$ to A. P_b sends $O_b = k_b + m_b$ to A.
3. **Broadcast phase**: A broadcasts $X = O_1 + ... + O_n$ to $P_1, ..., P_n$. P_a computes $m_b = X - m_a$ and P_b computes $m_a = X - m_b$.

In the description we assume that the setups for the protocols have already been performed. Typically one of the correspondents will anonymously provide the other players with the required data using an anonymous channel akin to [11, 6].

3.1 A Simple Protocol

We assume that during the execution of the protocol

- the number of players remains constant,
- no messages are lost (or arrive too late), and
- all players are honest.

In this situation two players P_a and P_b can communicate with each other in full-duplex using the classic DC-net principle. This approach is described in detail in Protocol 1. All players are initially provided with secret keys $k_1, ..., k_n$ that sum up to 0, so that the keys will vanish when all ciphertexts are aggregated. So if both P_a and P_b send during the same round, the result forwarded by the aggregator is the sum of their messages. By subtracting their own message from this result they can recover each other's message.

The previous system obviously fails, if only one user does not send anything or if his packet is lost. This seems to be a very strong restriction of the protocol.

3.2 A Packet-Loss Resilient Protocol

The problem of packet-loss is that if the aggregator does not receive all the packets, he can only make the sum X over a strict subset of $O_1, ..., O_n$. This means that the keys will not cancel and P_a and P_b cannot recover the messages like in the previous scenario.

To overcome this problem we can modify the previous protocol as shown in Protocol 2. During the initialization P_a and P_b are provided with the keys $k_1, ..., k_n$ of all players. Further the aggregator does not only broadcast the sum X, but also a list L informing which packets he received. Said list L informs the players P_a and P_b about which keys are included in the partial sum and since they know all the keys, they can subtract them from X and recover the messages.

This leaves us with the problem of users who deliberately send faulty packets to disrupt the communication. Such a case should be caught and the corresponding packets should be dropped. To identify

Protocol 2 A packet-loss resilient protocol.

1. **Setup**: $P_1, ..., P_n$ respectively have secret keys $k_1, ..., k_n$. P_a additionally has m_a and P_b additionally has m_b, and both know all $k_1, ..., k_n$.
2. **Collection phase**: Each P_i for $i \neq a, b$ sends $O_i = k_i$ to A. P_a sends $O_a = k_a + m_a$ to A. P_b sends $O_b = k_b + m_b$ to A. A receives O_i for $i \in L \subseteq \{1, ..., n\}$.
3. **Broadcast phase**: A broadcasts (L, X) to $P_1, ..., P_n$, where $X = \sum_{i \in L} O_i$. If $b \in L$ then P_a computes m_b; with $m_b = X - \sum_{i \in L} k_i - m_a$ in case $a \in L$ or otherwise with $m_b = X - \sum_{i \in L} k_i$. If $a \in L$ then P_b computes m_a; with $m_a = X - \sum_{i \in L} k_i - m_b$ in case $b \in L$ or otherwise with $m_a = X - \sum_{i \in L} k_i$.

such packets, we propose the following protocol.

3.3 A Protocol Resilient to Lost and Faulty Packets

The problem is that if a malicious player P_i sends a random value instead of k_i then P_a and P_b who expect k_i will not be able to properly extract the messages m_a and m_b from X anymore.

In order to protect against such malicious players it is obvious the aggregator must be able to distinguish if a received packet is correct. However the aggregator should not be able to distinguish which packets contain messages. Thus every player must include a proof that the submitted data is correct, and the aggregator must be able to verify this proof without gaining any other information from it. This means that a player $P_i \notin \{P_a, P_b\}$ must be able to prove that $O_i = k_i$, and P_a and P_b must keep the freedom to send $O_a = k_a + m_a$ and $O_b = k_b + m_b$.

An elegant way to achieve this is to bind each player P_i to his key k_i using a trapdoor commitment, where the secret trapdoor information α is only known to P_a and P_b. Then each player $P_i \notin \{P_a, P_b\}$ can only open the commitment to the value k_i, but P_a and P_b who know α can open their commitments to any value they like, that is to $k_a + m_a$ and $k_b + m_b$.

In our description we use Pedersen commitments [21] which are of the form $c = g^r h^k$, and we assume that the secret $\alpha = \log_g h$ is only known to P_a and P_b.

Protocol 3 A packet-loss resilient protocol with verification.

1. **Setup**: $P_1, ..., P_n$ respectively have secret value pairs $(k_1, r_1), ..., (k_n, r_n)$. P_a additionally has m_a and P_b additionally has m_b, and both know all $(k_1, r_1), ..., (k_n, r_n)$. A is provided with $c_1, ..., c_n$ where $c_i = g^{r_i} h^{k_i}$. Only P_a and P_b know $\log_g h$.
2. **Collection phase**: Each P_i sends (O_i, s_i) to A. Each P_i for $i \neq a, b$ uses $O_i = k_i$ and $s_i = r_i$. P_a uses $O_a = k_a + m_a$ and $s_i = r_i - m_a \cdot \log_g h$. P_b uses $O_b = k_b + m_b$ and $s_i = r_i - m_b \cdot \log_g h$. A receives (O_i, s_i) where additionally $g^{s_i} h^{O_i} = c_i$ holds for $i \in L \subseteq \{1, ..., n\}$.
3. **Broadcast phase**: A broadcasts (L, X) to $P_1, ..., P_n$, where $X = \sum_{i \in L} O_i$. If $b \in L$ then P_a computes m_b; with $m_b = X - \sum_{i \in L} k_i - m_a$ in case $a \in L$ or otherwise with $m_b = X - \sum_{i \in L} k_i$. If $a \in L$ then P_b computes m_a; with $m_a = X - \sum_{i \in L} k_i - m_b$ in case $b \in L$ or otherwise with $m_a = X - \sum_{i \in L} k_i$.

As shown in Protocol 3, during the setup the aggregator is provided with a commitment for each expected O_i, and each player P_i is provided with the corresponding secret pairs (k_i, r_i). Then, during the collection phase, each participant $P_i \notin \{P_a, P_b\}$ must send (k_i, r_i) to the aggregator, since without α he cannot find any other pair (k'_i, r'_i) that corresponds to the commitment. P_a and P_b can use α to compute valid pairs $(k_a + m_a, r'_a)$ and $(k_b + m_b, r'_b)$. The aggregator verifies for each received pair if it corresponds to the commitment and rejects pairs that do not. Thus only valid k_is are used to compute X.

This scenario is good if there is only one transmission round, but the anonymous sending of $c_1, ..., c_n$ to A is expensive and does not scale well to multiple rounds. Therefore we need a more efficient way to provide the aggregator A with means for verifying the data from the participants when there are multiple transmission rounds.

3.4 A Protocol Resilient to Packet-Loss and Malicious Players for Multiple Rounds

In order to extend the previous protocol to multiple transmission rounds, we propose to use Merkle

trees [20]. For each player P_i we use a Merkle tree T_i that allows to verify that a given commitment is valid for a given round. It is then not necessary anymore to provide the aggregator with a commitment for each round, but it is sufficient to provide the aggregator with the roots of the Merkle trees. As illustrated in Figure 3, such a Merkle tree T_i can be constructed from a sequence $(k_i^{(1)}, r_i^{(1)}), ..., (k_i^{(J)}, r_i^{(J)})$, which can be derived from a secret seed S_i.

So what changes compared to the preceding protocol is that the aggregator is provided with the roots of the Merkel trees instead of the commitments, and each player P_i is provided with a secret seed S_i that corresponds to a pseudorandom sequence. Then, during the transmission phase each player P_i does not only send $(k_i^{(j)}, r_i^{(j)})$ but $(k_i^{(j)}, r_i^{(j)}, Z_i^{(j)})$ where $Z_i^{(j)}$ is a proof the commitment corresponding to $(k_i^{(j)}, r_i^{(j)})$ is the right one for round j (i.e., that it is at position j in the sequence). The aggregator computes a commitment and verifies using $Z_i^{(j)}$ that it is correct. A detailed description of the protocol is shown in Protocol 4.

It is easy to see that this protocol satisfies the properties of correctness and privacy defined in Section 2.

3.5 Variants of the Protocol

In order to recover the messages in presence of packet-loss we proposed in Section 3.2 that the aggregator should send the list of packets that have been received along with the sum. This allows the receiver to directly compute the message. If one can assume that only a few (e.g., 1 or 2) packets are lost per round, one can also omit to send this list. The recipient can then still recover the message by trying all the possible combinations of missing packets. This way the packet length is reduced, however at the cost of a more expensive computation at the recipient.

Similarly one could also completely omit the cryptographic proof and go for a completely different mechanism. The players P_a and P_b could, upon detection of problems, use the anonymous channel from the setup and ask the aggregator to publish all the packets he received during a problematic round. Since P_a and P_b know all the keys, they would directly distinguish who sent a faulty packet, and they could anonymously ban those players from the group. This optimistic approach would lead to shorter packets, but as the latency of the anonymous channel is expected to be high, the stream would be interrupted for a non negligible amount of time.

Protocol 4 A packet-loss resilient protocol with verification and multiple rounds.

1. **Setup**: For each round $j \in \{1, ..., J\}$, $P_1, ..., P_n$ respectively have secret value pairs $(k_1^{(j)}, r_1^{(j)}), ..., (k_n^{(j)}, r_n^{(j)})$. P_a additionally has $m_a^{(j)}$ and P_b additionally has $m_b^{(j)}$, and both know all $(k_1^{(j)}, r_1^{(j)}), ..., (k_n^{(j)}, r_n^{(j)})$. A is provided with $R_1, ..., R_n$ the roots of a merkle trees $T_1, ..., T_2$ constructed from $(c_1^{(1)}, ..., c_1^{(J)}), ..., (c_n^{(1)}, ..., c_n^{(J)})$ where $c_i^{(j)} = g^{r_i^{(j)}} h^{k_i^{(j)}}$. Only P_a and P_b know $\log_g h$.

2. **Collection phase** (round j): Each P_i sends $(O_i^{(j)}, s_i^{(j)}, z_i^{(j)})$ to A. Each P_i for $i \neq a, b$ uses $O_i^{(j)} = k_i^{(j)}$ and $s_i^{(j)} = r_i^{(j)}$. P_a uses $O_a^{(j)} = k_a^{(j)} + m_a^{(j)}$ and $s_i^{(j)} = r_i^{(j)} - m_a^{(j)} \cdot \log_g h$. P_b uses $O_b^{(j)} = k_b^{(j)} + m_b^{(j)}$ and $s_i^{(j)} = r_i^{(j)} - m_b^{(j)} \cdot \log_g h$. Further, $z_i^{(j)}$ is a proof that $c_i^{(j)} = g^{r_i^{(j)}} h^{k_i^{(j)}}$ is in the Merkle tree T_i at position j. A receives $(O_i^{(j)}, s_i^{(j)}, z_i^{(j)})$ where $z_i^{(j)}$ proves that $c_i^{(j)} = g^{s_i^{(j)}} h^{O_i^{(j)}}$ is at position j in T_i, for $i \in L^{(j)} \subseteq \{1, ..., n\}$.

3. **Broadcast phase** (round j): A broadcasts $(L^{(j)}, X^{(j)})$ to $P_1, ..., P_n$, where $X^{(j)} = \sum_{i \in L^{(j)}} O_i^{(j)}$. If $b \in L^{(j)}$ then P_a computes $m_b^{(j)}$; with $m_b^{(j)} = X^{(j)} - \sum_{i \in L^{(j)}} k_i^{(j)} - m_a^{(j)}$ in case $a \in L^{(j)}$ or otherwise with $m_b^{(j)} = X^{(j)} - \sum_{i \in L^{(j)}} k_i^{(j)}$. If $a \in L^{(j)}$ then P_b computes $m_a^{(j)}$; with $m_a^{(j)} = X^{(j)} - \sum_{i \in L^{(j)}} k_i^{(j)} - m_b^{(j)}$ in case $b \in L^{(j)}$ or otherwise with $m_a^{(j)} = X^{(j)} - \sum_{i \in L^{(j)}} k_i^{(j)}$.

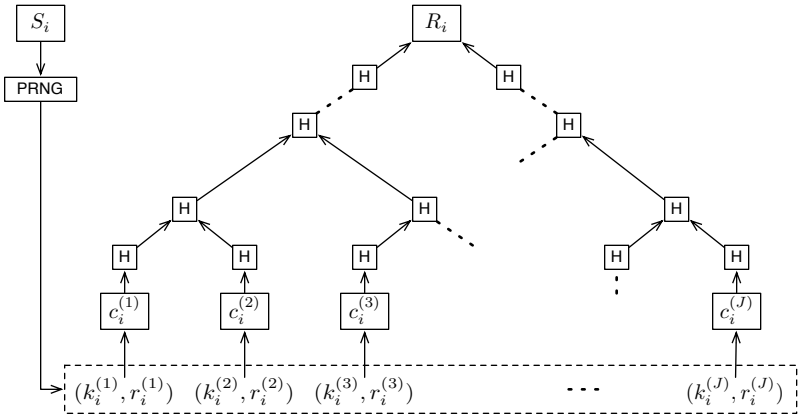

Figure 3: A Merkle tree T_i for the sequence of commitments $c_i^{(1)}, ..., c_i^{(J)}$ of player P_i. The sequence of the secret keys $(k_i^{(1)}, r_i^{(1)}), ..., (k_i^{(J)}, r_i^{(J)})$ can be pseudo-randomly generated from a secret seed S. The player P_i can prove that the commitment $c_i^{(j)}$ is at the position j in the sequence.

4 Practical considerations

In this section we discuss some aspects to consider in a real implementation of the protocol.

Channel Setup Concerning the setup of the channel, we assumed so far that the initialization is performed anonymously by one of the correspondents P_a or P_b. This correspondent will provide all other players and the aggregator with the the required data via an anonymous communication channel.

One way to implement such an anonymous channel is to use a DC-net. However, such a DC-net must then be run periodically, since in general it is not known in advance when a correspondent will want to talk with another. The higher the frequency with which such a DC-net is run, the better the user experience. But as each run consumes bandwidth, one does not want to do this more often than necessary either. So there is a tradeoff to be made between bandwidth and user experience with this approach.

Another way to implement an anonymous channel is to use a relay based approach like onion routing (e.g. TOR). Here the problem is that the overall security provided by the system is only as strong as the weakest link in the chain. The use of such a relay based approach would weaken the overall security of the system.

Channel Termination In the protocol of Section 3.4 the number of rounds (and thus the length of the call) is fixed during the setup of the channel. If a call ends earlier the correspondents can actively terminate the call by notifying the other players via the same anonymous channel that they used to do the initialization.

Load Distribution with P2P Especially for the aggregator the computational costs and the bandwidth requirements and can rise to non-negligible levels, since they are proportional to the number of players. For instance if all packets are around 100 bytes and if 100 players send 50 packet per second, the aggregator must aggregate 5000 packets per second and has a corresponding incoming and outgoing traffic of 4 Mbps. Each participant would have an incoming an outgoing traffic of 40 kbps.

In a P2P system one is not obliged to have only one aggregator as illustrated in Figure 4a, but the players can distribute this load between all of them by successively have each one of them play the role of the aggregator in a round robin fashion as illustrated in Figure 4b. This way the load is more evenly distributed and for the same setup as in the preceding example each player would have of around 80 kbps of ignoring and outgoing traffic.

Synchronization All players should send their packets such that they arrive at the aggregator practically at then same time, in order to minimize the overall latency of a transmission. The aggregator will only wait for a certain period of time before aggregating the received data and sending the result to the players. It is therefore important that the clocks of

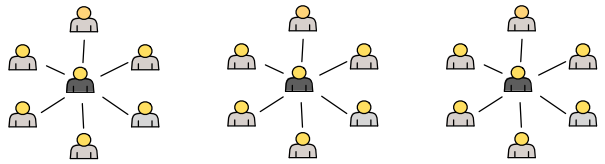

(a) Implementation with a dedicated aggregator: In each round the same aggregator processes the data from all the participants.

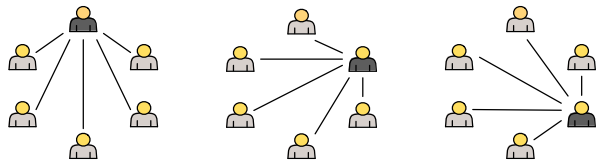

(b) Implementation as a distributed P2P system: In each round another party plays the role of the aggregator. Advantageously this also mitigates the damage that dishonest aggregator can cause.

Figure 4: Centralized and distributed systems.

the players are properly synchronized.

Cryptography The cryptographic assumption for the commitments and the hashtables only needs to hold for a short time. It is therefore possible to use a lower security parameter than for digital signatures that have to be secure for decades.

5 Performance

In this section we consider the latency, the bandwidth and the computing complexity for some setups.

Latency Latency of packets in computer networks is often assumed to follow a log-normal distribution, e.g. in [17, 22]. This distribution defined by

$$Pr(t = x) = \frac{1}{\sqrt{2\pi s^2}} \cdot \exp\left(\frac{(\ln x - u)^2}{2s^2}\right)$$

has a characteristic heavy tail, as shown shown in Figure 5a. If we assume an average $u = 0.97$ and a standard error $s = 0.06$, then the average time the aggregator has to wait until all n independent ciphertexts have arrived is given by

$$Pr(t < x)^n = \left(\int_{l=-\infty}^{x} Pr(t = l)\right)^n.$$

where n is the number of players. As shown in Figure 5b the cumulated latency increases with the number of players. In our case, we see that for $n = 100$ players we already have a latency increase of more than 30ms.

Packet Loss Packet loss typically occurs in bursts and can be modeled using the well known Gilbert–Elliott (GE) channel [14, 10]. We estimate the number of rounds during which no packet is lost on its way to the aggregator, based on the probability p that a packet is lost. The probability that a packet is not lost is then $1 - p$, and the probability that no packet is lost is then

$$q = (1 - p)^n.$$

Figure 6 illustrates the number of rounds during which at least one packet does not make it to the aggregator for various values of p.

Bandwidth During one transmission round each of the n players sends a packet to the aggregator, and the aggregator sends a packet back to each player. The total number of packets per second $b(n)$ is thus proportional to the number of players n. That is

$$b(n) = 2 \cdot \frac{n - 1}{f},$$

where f is the number of rounds per second.

The load of the aggregator increases with the number of players.

In a P2P scenario where all players successively play the role of the aggregator, the bandwidth usage is distributed evenly amongst all players. Each player will perform like a normal player for n-1 rounds, and in one out of n rounds he will not have to send anything, but he will have to broadcast the aggregated data to the $n - 1$ other players. The bandwidth per player $p(n)$ is shown in Figure 7a. It can be computed with

$$p(n) = \frac{b}{n} = \frac{n - 1}{n} \cdot \frac{2}{f} \sim \frac{2}{f}.$$

Packet size Packets in our protocol are composed of two parts, the audio payload on one hand and the cryptographic overhead on the other hand.

The amount of audio data depends on the frequency of the packets, on the quality (sampling frequency, compression rate) of the sound and the number of sound channels (e.g., mono or stereo). For voice transmission in mono this could be 50 packets with 60 bytes per second per packet, but for high-end music in stereo it will be significantly more.

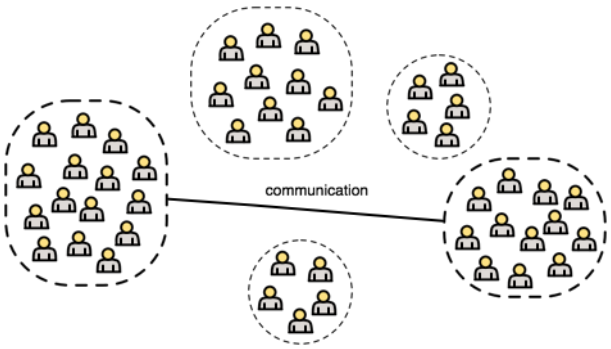

Figure 8: In an untraceable communication system with groups, there is a whole group of potential correspondents on either side of the line. Nobody except the real correspondents – neither an external observer, nor any other group member – can distinguish who is communicating with whom.

The amount of cryptographic data depends on the strength of the cipher that is used. Since the cryptographic assumption only has to hold during the communication one can use lightweight cryptographic primitives.

6 Future Work

There are basically two directions for future work, the improvement of the protocol itself and the building of larger systems composed of multiple groups.

Detection of Malicious Aggregators In this paper we assume that the aggregator is honest but curious. This means that he will not drop packets, nor omit to send packets. While he cannot send wrong results and remain undetected, he can just disrupt a transmission round by just dropping data. A more powerful malicious aggregator could however just ignore some of the packets he receives, or he could deliberately not broadcast the aggregated result to anybody. In a P2P setting the effect of such an aggregator can be mitigated using the rotation principle proposed in section X, but ideally one would like to detect and to ban such aggregators from the group.

Larger Systems with Multiple Groups Protocols based on DC-net do not scale to a very large number of participants, as the bandwidth and the computational power used by the aggregator are proportional to the number of participants. So the idea which was already proposed in [15] is to realize systems composed of many small groups, as illustrated in Figure 8. Only the correspondents will know that they are communicating, all other players or observers cannot distinguish who is communicating.

For example one could have on one side a group of 500 politicians and on the other side a group of 500 journalists. When a politician then talks to a journalist, it would only be possible to see that one of the 500 politicians is talking to one of the 500 journalists, but it would not be possible for anybody to distinguish which politician is talking to which journalist. As there would be $500 \cdot 500 = 250000$ different possibilities, such a system would offer a fairly good protection.

There are many open questions, such as: How can we locate a given participant within the system, if we do not know in which group he is? How can we handle participants joining and leaving the system? How can we ban malicious participants from the entire system?

7 Related Work

The Dining Cryptographers protocol was proposed in [4] and further studied in [3, 2, 23, 24]. A first system composed of multiple DC-nets was proposed in [15].

Computationally secure DC-net protocols with zero-knowledge verification of the data have been proposed in [16] and further studied in [11, 25, 6, 12, 7].

Recent group based communication systems include [7, 18, 19, 8]. There have also been TOR [9] extensions for VoIP [13] and for group communication [26].

8 Concluding Remarks

Starting from the classic DC-net paradigm we derived a protocol for untraceable VoIP telephony that is resilient against packet-loss and faulty packets. It enables two players within a larger group to communicate with VoIP without anybody else being able to distinguish that they are communicating. Further we discussed practical issues and showed how to distribute the load in a P2P network.

We consider this work a first step towards larger systems composed of multiple groups so that can scale to a larger number of participants.

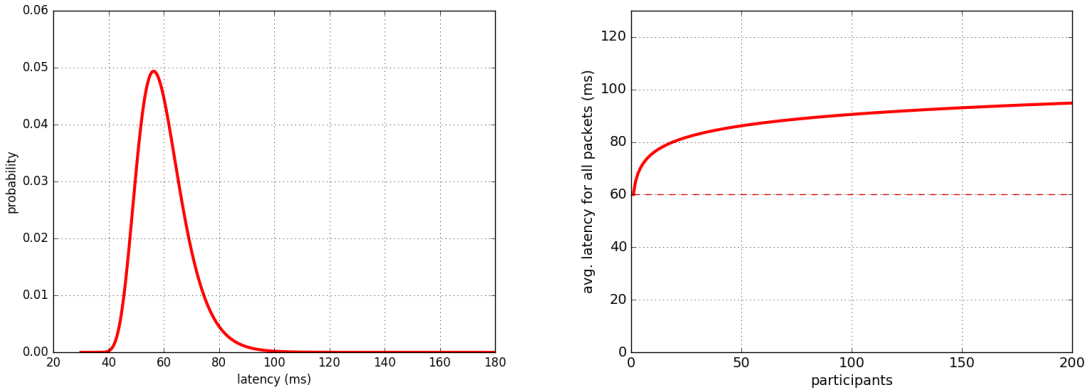

(a) Latency of a packet.

(b) Average latency of the last packet.

Figure 5: Latency of arrival of the packets at the aggregator.

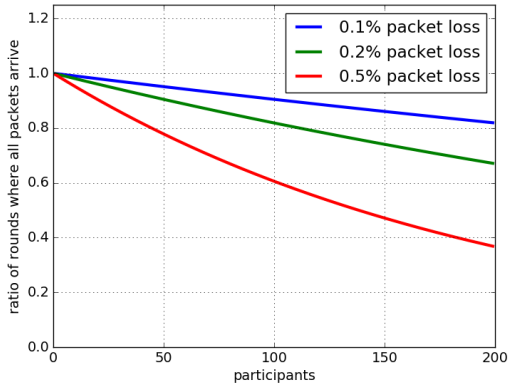

Figure 6: Ratio of rounds in which no packet is lost on its way to the aggregator.

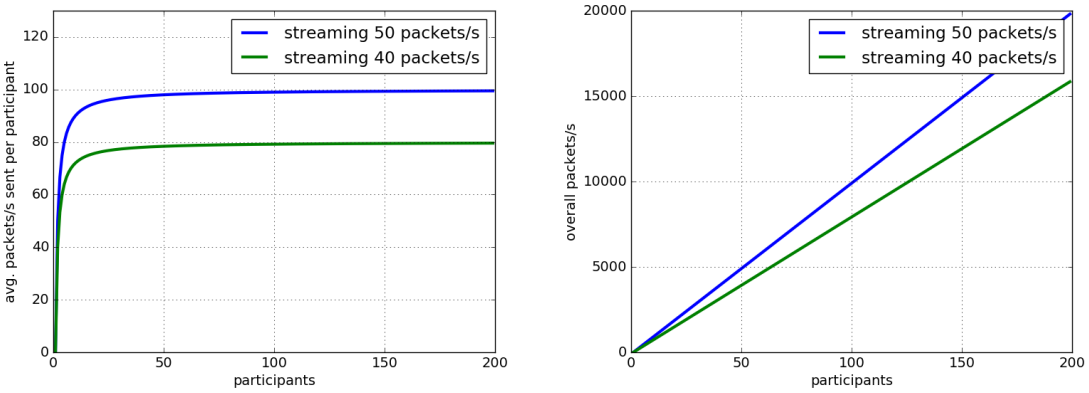

(a) Packets sent per participant.

(b) Total number of packets sent.

Figure 7: Bandwidth

References

[1] One-way transmission time. Recommendation G.114, International Telecommunication Union, May 2003.

[2] Waidner M. Baum-Waidner B., Pfitzmann B. Unconditional byzantine agreement with good majority. 1991.

[3] J. Bos and B. den Boer. Detection of disrupters in the DC protocol. *Advances in Cryptology–EUROCRYPT'89*, pages 320–327, 1989.

[4] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1(1):65–75, 1988.

[5] Cisco. *Voice Over IP - Per Call Bandwidth Consumption*, 2016. Available at http://www.cisco.com/c/en/us/support/docs/voice/voice-quality/7934-bwidth-consume.html (accessed 10/2016).

[6] H. Corrigan-Gibbs, D. I. Wolinsky, and B. Ford. Proactively accountable anonymous messaging in verdict. In *USENIX Security*, 2013.

[7] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE Symposium on Security and Privacy*, pages 321–338. IEEE, 2015.

[8] George Danezis, Claudia Diaz, Carmela Troncoso, and Ben Laurie. Drac: An architecture for anonymous low-volume communications. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 202–219. Springer, 2010.

[9] R. Dingledine, N. Mathewson, and P. Syverson. Tor: the second-generation onion router. *Proceedings of the 13th conference on USENIX Security Symposium-Volume 13 table of contents*, pages 21–21, 2004.

[10] Edwin O Elliott. Estimates of error rates for codes on burst-noise channels. *Bell system technical journal*, 42(5):1977–1997, 1963.

[11] C. Franck. New Directions for Dining Cryptographers. Master's thesis, University of Luxembourg, Luxembourg, 2008.

[12] Christian Franck. Dining cryptographers with 0.924 verifiable collision resolution. *Annales UMCS, Informatica*, 14(1):49–59, October 2014.

[13] Van Gegel. Onionphone v0. 1a: Voip add-on for torchat.

[14] Edgar N Gilbert. Capacity of a burst-noise channel. *Bell system technical journal*, 39(5):1253–1265, 1960.

[15] S. Goel, M. Robson, M. Polte, and E. Sirer. Herbivore: A Scalable and Efficient Protocol for Anonymous Communication. 2003.

[16] P. Golle and A. Juels. Dining Cryptographers Revisited. *Advances in cryptology-EUROCRYPT 2004: International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004: Proceedings*, 2004.

[17] Bernardo A Huberman and Rajan M Lukose. Social dilemmas and internet congestion. *Science*, 277(5325):535–537, 1997.

[18] Albert Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle. *Proceedings on Privacy Enhancing Technologies*, 2016(2):115–134, 2015.

[19] Stevens Le Blond, David Choffnes, William Caldwell, Peter Druschel, and Nicholas Merritt. Herd: A scalable, traffic analysis resistant anonymity network for voip systems. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 639–652. ACM, 2015.

[20] Ralph C Merkle. A certified digital signature. In *Conference on the Theory and Application of Cryptology*, pages 218–238. Springer, 1989.

[21] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *CRYPTO*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 1991.

[22] U. Sorger and Z. Suchanecki. Delays of interconnected flows. *International Journal of Theoretical Physics*, 50(7):2242–2251, 2011.

[23] M. Waidner. Unconditional Sender and Recipient Untraceability in spite of Active Attacks. *Lecture Notes in Computer Science*, 434:302, 1990.

[24] M. Waidner and B. Pfitzmann. The Dining Cryptographers in the Disco: Unconditional Sender and Recipient Untraceability with Computationally Secure Serviceability. *Advances in Cryptology - EUROCRYPT*, 89:10–13, 1989.

[25] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In *Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, pages 179–182, 2012.

[26] Lei Yang and Fengjun Li. mtor: a multipath tor routing beyond bandwidth throttling. In *Communications and Network Security (CNS), 2015 IEEE Conference on*, pages 479–487. IEEE, 2015.